

A game theoretic approach to a peer-to-peer cloud storage model

Fabio Fagnani¹, Barbara Franci², and Ennio Grasso³

Abstract—Classical cloud storage based on external data providers has been recognized to suffer from a number of drawbacks. This is due to its inherent centralized architecture which makes it vulnerable to external attacks, malware, technical failures, as well to the large premium charged for business purposes. In this paper, we propose an alternative distributed peer-to-peer cloud storage model which is based on the observation that the users themselves often have available storage capabilities to be offered in principle to other users. Our set-up is that of a network of users connected through a graph, each of them being at the same time a source of data to be stored externally and a possible storage resource. We cast the peer-to-peer storage model to a Potential Game and we propose an original decentralized algorithm which makes units interact, cooperate, and store a complete back up of their data on their connected neighbors. We present theoretical results on the algorithm as well a good number of simulations which validate our approach.

I. INTRODUCTION

Cloud storage on the Internet has come to rely almost exclusively on data providers serving as trusted third parties to transfer and store the data. While the system works well enough in most cases, it still suffers from the inherent weaknesses of the trust-based model. The traditional cloud is open to a variety of security threats, including man-in-the-middle attacks, and malware that expose sensitive and private consumer and corporate data. Furthermore, current cloud storage applications are charging large premiums on data storage facilities to business clients. Moreover, these cloud storage providers may have technical failures that can cause data breaches and unavailability, much to the distress of the users and applications that depend on them.

To address the aforementioned shortcomings, a decentralized peer-to-peer cloud storage model would be the right answer. On the wake of the successful peer-to-peer file sharing model of applications like BitTorrent and its lookalike, the same philosophy may well be leveraged on a different but very similar application service like storage. Indeed a slew of fledgling and somehow successful startups are entering in this market niche. Among the most noteworthy examples are:

- Storj: www.youtube.com/channel/UC-cTEqWwZV5R1-h0RZsp2Qw,
- BitTorrent Sync: www.getsync.com/,
- Ethos: www.youtube.com/watch?v=qUftGCQ5dpo,

- SpaceMonkeys: www.spacemonkey.com/.

Clearly, a completely decentralized peer-to-peer model must account for some challenging technical difficulties that are absent in a centralized cloud model. Firstly, security and privacy must be carefully implemented by ensuring end-to-end encryption resistant to attackers. In addition, the model must account for the latency, performance, and downtime of average user devices.

Albeit the above technical issues are challenging, they can be addressed with the right tools and architectures available at current state of the art technology and will not be considered in this paper. What remains an open question up to now is how to endow the system with the right incentives for the end users to collaborate and share their storage commodity with each other. We believe that the answer to that question comes from the formal framework of game-theory which provides the mathematical tools to ensure successful cooperation among users / players. There are (at least) two ways in which game-theory, with its plenty folk-theorems, can be applied to the real world. One is as a tool to study an ongoing phenomenon. This is the typical setting of social and psychological sciences. A second more engineering approach is to leverage game-theory to design specific mechanisms, i.e. set rules of the game that will bring the interaction to the most desired outcome, which is basically the maximum global welfare, or Pareto dominated equilibrium of the game. This work follows this second approach.

In this paper, we consider a network of units (PC's but possibly also smartphones or other devices possessing storage capabilities) which need to store externally a back up of their data and, at the same time, can offer space available to store data of other connected units. In this set up, we cast the peer-to-peer storage model to an allocation Potential Game and we propose an original decentralized algorithm which make units interact, cooperate, and store a complete back up of their data on their connected neighbors.

Units are assumed to be connected through a network and, autonomously, at random time, to activate and allocate or move their data pieces among the neighboring units. Formally, each unit has a utility function which gives a value to their neighbors on the basis of their reliability, their current congestion (resources have bounded storage capabilities), and the amount of data the unit has already stored in them. Following classical evolutionary game theory [6], we propose an algorithm based on a noisy best response action: each time a unit activates, it decides the neighbor to use on the basis of a Gibbs probability distribution having its peak on the maxima of the utility function.

¹Fabio Fagnani is with the Department of Mathematical Sciences, Politecnico di Torino, fabio.fagnani@polito.it

²Barbara Franci is with the Department of Mathematical Sciences, Politecnico di Torino, barbara.franci@polito.it

³Ennio Grasso is with Swarm Lab - Joint Open Lab TelecomItalia, Torino ennio.grasso@telecomitalia.it

In the remaining part of this section, we formally define the storage allocation problem and we show its equivalence with classical matching problem on a graph. This allows us to use celebrated Hall's theorem and give a necessary and sufficient condition for the allocation problem to be solvable. Section II is devoted to cast the problem to a potential game theoretic framework [2] and to propose a distributed algorithm which is an instance of a noisy best response dynamics [6]. We claim a fundamental result, which will be proven in a forthcoming paper, which says that in the double limit when time goes to infinity and the noise parameter goes to 0, the algorithm converges to a Nash equilibrium which is, in particular, a global maximum of the potential function. This guarantees that the solution will indeed be close to the global welfare of the community. Finally, Section III is devoted to the presentation of an extensive set of simulations. A conclusions section ends the paper.

A. Related Work

Though allocation games have been considered before in the literature, they all substantially differ from the model we propose and study in this paper. In [4] the authors consider an allocation problem casted to a pure congestion game. Utility functions of units measure the congestion of a resource simply as a function of the number of units currently using it, but they do not impose any strict storage limitation. The algorithm they propose is a classical best response algorithm and is shown to achieve Nash equilibrium. Our model differs considerably as we also consider reliability of the resources and data fragmentation in the utility functions and, moreover, we impose strict storage limitations. A crucial consequence of this is that classical best response algorithms would not work in our case: Example 3 shows a situation where such an algorithm would halt before allocation is completed. Allocation games are also considered in [1] where, however, the proposed algorithm units are not interacting through a graph but rather through a device that acts like a leader selecting which resources can be used. A related context where congestion games have been used is that of networking routing [3].

At a broader level, the noisy best response algorithm we propose in this paper fits in the so-called evolutionary game theory [6] which has already been extensively used in studying other networking problems [5].

B. The model

Consider a set \mathcal{X} of units which play the double role of users who have to allocate externally a back up of their data, as well resources where data from other units can be allocated. Generically, an element of \mathcal{X} will be called a unit, while the terms user and resource will be used when the unit is considered in the two possible roles of, respectively, a source or a recipient of data. We assume units to be connected through a directed graph $\mathcal{G} = (\mathcal{X}, \mathcal{E})$ where a link $(x, y) \in \mathcal{E}$ means that unit x is allowed to store data in unit y . We denote by

$$N_x := \{y \in \mathcal{X} \mid (x, y) \in \mathcal{E}\}, \quad N_y^- := \{x \in \mathcal{X} \mid (x, y) \in \mathcal{E}\}$$

respectively, the out- and the in-neighborhood of a node. Note the important different interpretation in our context: N_x represents the set of resources available to unit x while N_y^- is the set of units having access to resource y . If $D \subseteq \mathcal{X}$, we put $N(D) = \cup_{x \in D} N_x$ and $N^-(D) = \cup_{x \in D} N_x^-$.

We imagine the data possessed by the units to be quantized atoms of the same size. Each unit x is characterized by two non negative integers:

- α_x is the number of data atoms that unit x needs to back up into his neighbors,
- β_x is the number of data atoms that unit x can accept and store from his neighbors.

The numbers $\{\alpha_x\}$ and $\{\beta_x\}$ will be assembled into two vectors denoted, respectively, α and β . We also define

$$\mathcal{A}_x = \{(x, a) \mid a \in \{1, \dots, \alpha_x\}\}, \quad \mathcal{A} = \bigcup_{x \in \mathcal{X}} \mathcal{A}_x$$

Given the triple $(\mathcal{G}, \alpha, \beta)$, we define an allocation as any map $Q : \mathcal{A} \rightarrow \mathcal{X}$ satisfying the properties expressed below.

- (C1) **Graph constraint** $Q(x, a) \in N_x$ for all $x \in \mathcal{X}$ and $a \in \{1, \dots, \alpha_x\}$;
- (C2) **Storage limitation** For every $y \in \mathcal{X}$,

$$|Q^{-1}(y)| \leq \beta_y$$

The fact that $Q(x, a) = y$ means that user x has allocated the data atom a into resource y .

We will say that the allocation problem is solvable if an allocation Q exists. We denote by \mathcal{Q} the set of allocations. We will also need to consider partial allocations, namely maps $Q : D \rightarrow \mathcal{X}$ where $D \subseteq \mathcal{A}$ satisfying, where defined, conditions (C1) and (C2). We denote by \mathcal{Q}_p the set of partial allocations.

In the following section we study the conditions under which the allocation problem is solvable, namely conditions under which \mathcal{Q} is non empty.

C. The allocation problem as a matching problem

Define

$$\mathcal{B}_y = \{(y, b) \mid b \in \{1, \dots, \beta_y\}\}, \quad \mathcal{B} = \bigcup_{y \in \mathcal{X}} \mathcal{B}_y$$

Consider now the bipartite graph $\mathcal{P} = (\mathcal{A} \times \mathcal{B}, \mathcal{E}_{\mathcal{P}})$ where $((x, a), (y, b)) \in \mathcal{E}_{\mathcal{P}}$ iff $(x, y) \in \mathcal{E}$. An allocation naturally induces a matching on \mathcal{P} which is complete on \mathcal{A} . To this aim, notice that, from $Q \in \mathcal{Q}$ and using condition (C2), we can construct an injective mapping $\tilde{Q} : \mathcal{A} \rightarrow \mathcal{B}$ such that $\tilde{Q}(x, a) = (Q(x, a), b)$ for every $x \in \mathcal{X}$ and for all a . We then define

$$\mathcal{M} := \bigcup_{x \in \mathcal{X}} \{((x, a), (y, b)) \in \mathcal{A} \times \mathcal{B} \mid (y, b) = \tilde{Q}(x, a)\}$$

It is clear that this procedure can be inverted and that from any matching of \mathcal{M} complete on \mathcal{A} we can associate an allocation for $(\mathcal{G}, \alpha, \beta)$. This equivalence allows to use classical results like the Hall's marriage theorem to characterize the existence of allocations. Precisely we have the following result

Theorem 1: Given $(\mathcal{G}, \alpha, \beta)$, there exists an allocation iff the following condition is satisfied:

$$\sum_{x \in D} \alpha_x \leq \sum_{y \in N(D)} \beta_y \quad \forall D \subseteq \mathcal{X} \quad (1)$$

Proof: By Hall's theorem, the existence of a matching in \mathcal{P} complete on A is equivalent to the condition

$$|A| \leq |N^{\mathcal{P}}(A)| \quad \forall A \subseteq \mathcal{A} \quad (2)$$

where $N^{\mathcal{P}}(A) \subseteq \mathcal{B}$ is the out-neighborhood of A in \mathcal{P} . Given $A \subseteq \mathcal{A}$ let \bar{A} be the union of those \mathcal{A}_x 's for which $\mathcal{A}_x \cap A \neq \emptyset$. By the way the bipartite graph \mathcal{P} has been defined, it follows that $N^{\mathcal{P}}(A) = N^{\mathcal{P}}(\bar{A})$, so that it is sufficient to restrict condition (2) to subsets A such that $\mathcal{A}_x \cap A \neq \emptyset$ yield $\mathcal{A}_x \subseteq A$. Given such an A , if we consider $D = \{x \mid \mathcal{A}_x \subseteq A\}$, we immediately obtain that (1) coincides with (2). ■

In general, it is not necessary to check the validity of (1) for every subset D . We say that $D \subseteq \mathcal{X}$ is maximal if for any $D' \supsetneq D$, it holds $N(D') \supsetneq N(D)$. We say that $D_1, D_2 \subseteq \mathcal{X}$ are independent if $N(D_1) \cap N(D_2) = \emptyset$ and $D \subseteq \mathcal{X}$ is called irreducible if it can not be decomposed into the union of two non empty independent subsets. Clearly, it is sufficient to verify (1) for the subclass of maximal irreducible subsets.

Example 1: If \mathcal{G} is complete, we have that $N(\{x\}) = \mathcal{X} \setminus \{x\}$ while $N(D) = \mathcal{X}$ for all D such that $|D| \geq 2$. Hence, the only maximal irreducible subsets are the singletons $\{x\}$ and the set \mathcal{X} . Condition (1) in this case reduces to

$$\alpha_x \leq \sum_{y \neq x} \beta_y, \forall x \in \mathcal{X} \quad \sum_{x \in \mathcal{X}} \alpha_x \leq \sum_{y \in \mathcal{X}} \beta_y \quad (3)$$

In general, the class of maximal irreducible subsets can be large and grow more than linearly in the size of \mathcal{X} , as the following example shows.

Example 2: If $\mathcal{G} = (\mathcal{X}, \mathcal{E})$ is a line graph ($\mathcal{X} = \{1, 2, \dots, n\}$ and $\mathcal{E} = \{(i, i+1), i = 1, \dots, n-1\}$) it can be checked that the maximal irreducible subsets are those of the form $\{i, i+2, \dots, i+2s\}$.

In case when $\alpha_x = a$ and $\beta_y = b$ are both constant, something more can be said.

Proposition 2: Given $(\mathcal{G}, \alpha, \beta)$, where \mathcal{G} is regular, $\alpha_x = a$ and $\beta_y = b$ for all $x, y \in \mathcal{X}$, there exists an allocation iff $a \leq b$.

Proof: By Theorem 1, we simply have to show that $|D| \leq |N(D)|$ for every subset $D \subseteq \mathcal{X}$. Let \mathcal{E}_D be the set of edges having one of the nodes in D . If d is the degree of the nodes in \mathcal{G} , we have that

$$d|D| = |\mathcal{E}_D| \geq d|N(D)|$$

■

From the practical point of view, the equivalence of our problem with a classical matching problem, is, however, of little utility, as the number of nodes of \mathcal{P} is of the size $\sum \alpha_x + \sum \beta_y$ which will in general be very large.

Moreover, in case allocations exist, we want to be able to construct one in a distributed way without the need of any supervision. The algorithm must be iterative in order to cope with possible time modifications of the units, of their interconnection and of their data storage needs and

limitations. Also the possibility that units leave and enter the community must be considered.

Also we want to have the possibility to find solution possessing certain extra features:

- **(Reliability)** resources will often have different reliability properties and we want to give preference, in the allocation, to more reliable resources;
- **(Congestion)** equally reliable resources should be equally used, avoiding congestion phenomena;
- **(Aggregation)** users prefer to use as few resources as possible to allocate their back up data.

The reason for this last feature comes from the fact that an exceeding fragmentation of the back up data will cause a blow up in the number of communications among the units both in the storage and recover phases. This feature should be considered against another feature which in this paper is not going to be addressed, which is that of diversification of back ups: in real applications units will need to back up multiple copies of their data in order to cope with security and possible failure phenomena. In that case, these multiple copies will need to be stored in different units. This issue will be analyzed in a subsequent paper.

The above desired features may be contradictory in general and we want to have tunable parameters to make the algorithm converge towards a desired compromised solution.

The proposed algorithm will be fully distributed: units will iteratively allocate and move their data among the neighbors on the basis of the space available and trying to maximize a utility function. There will be an underlying game theoretic structure inspired by the desired features described above. Our algorithm will be analyzed with the techniques of evolutionary game theory and it will be shown to yield a reversible Markov process converging to a Nash equilibrium of the game.

II. THE GAME THEORETIC SET-UP AND THE ALGORITHM

Given a partial allocation $Q \in \mathcal{Q}_p$, consider the matrix $W(Q) \in \mathbb{N}^{\mathcal{X} \times \mathcal{X}}$ where $W(Q)_{xy}$ is the number of atomic data that x has copied inside y under the allocation Q , namely,

$$W(Q)_{xy} := |Q^{-1}(y) \cap \mathcal{A}_x| \quad (4)$$

Clearly $W = W(Q)$ satisfies the following conditions

- (P1) $W_{xy} \geq 0$ for all x, y and $W_{xy} = 0$ whenever $(x, y) \notin \mathcal{E}$.
- (P2) $W^x := \sum_{y \in \mathcal{X}} W_{xy} \leq \alpha_x$ for all $x \in \mathcal{X}$.
- (P3) $W_y := \sum_{x \in \mathcal{X}} W_{xy} \leq \beta_y$ for all $y \in \mathcal{X}$.

It is immediate to see that, conversely, if there exists W satisfying these properties (such a W is called a partial allocation state), then, from it, we can construct a partial allocation Q such that $W = W(Q)$. Clearly, under this correspondence, we have that $Q \in \mathcal{Q}$ iff W satisfies (P2) with equality for all $x \in \mathcal{X}$. In this case W is called an allocation state. The set of partial allocation states and the set of allocation states are denoted, respectively, with the symbols \mathcal{W}_p and \mathcal{W} .

It is clear that two partial allocations Q^1 and Q^2 such that $W(Q^1) = W(Q^2)$, only differ for a permutation of the data

atoms of the various units and for many purposes can be considered as equivalent. All the quantity of interest for the game theoretic setting will be defined at the level of W .

We are now ready to define the game theoretic model. We first define utilities: under a (possibly partial) allocation state W , the utility of a unit x in using resource y is given by

$$f_{xy}(W) := \lambda_y - k_c W_y / \beta_y + k_a W_{xy} \quad (5)$$

The first term λ_y encodes possible reliability differences among resources, the second term is instead a congestion term which takes into consideration the level of use of the resource, and, finally, the third term depends on both x and y and pushes a unit to allocate in those resources where it has already allocated. k_c, k_a are two non-negative parameters to tune the effect of the congestion and of the aggregation terms, respectively.

The choice of this particular utility function has been made on the basis of simplicity considerations (notice that the state W enters linearly in it) and on the fact that, as exploited below, this leads to a potential game. In principle, different terms in the utility function can be introduced in order to make units to take into considerations other desired features (e.g multiple back up).

Define

$$\Psi(W) := \sum_{y \in \mathcal{X}} \sum_{s=0}^{W_y} [\lambda_y - k_c s / \beta_y] + k_a \sum_{x,y \in \mathcal{X}} \sum_{s=0}^{W_{xy}} s \quad (6)$$

and notice that if $W, W' \in \mathcal{W}$ are such that $W' = W - e_{x\bar{y}} + e_{x\bar{y}'}$, then,

$$\begin{aligned} \Psi(W') - \Psi(W) &= (\lambda_{\bar{y}'} - k_c W'_{\bar{y}'} / \beta_{\bar{y}'} + k_a W'_{x\bar{y}'}) - (\lambda_{\bar{y}} - k_c W_{\bar{y}} / \beta_{\bar{y}} + k_a W_{x\bar{y}}) \\ &= f_{x\bar{y}'}(W') - f_{x\bar{y}}(W) \end{aligned} \quad (7)$$

In other terms, under the state allocation W , when user x moves a data atom from \bar{y} to \bar{y}' , it experiences a variation in utility given by $\Psi(W') - \Psi(W)$. Ψ is called a potential of the game. Given $W \in \mathcal{W}_p$ and $x \in \mathcal{X}$, put

$$\mathcal{X}^x(W) := \{y \in N_x \mid W_y < \beta_y\}$$

the set of resources still available for x under the current allocation state W .

An allocation state $W \in \mathcal{W}$ (and also any $Q \in \mathcal{Q}$ such that $W(Q) = W$) is called a Nash equilibrium if, for every $x \in \mathcal{X}$, for every $y \in N_x$ such that $W_{xy} > 0$, for every $y' \in \mathcal{X}^x(W)$, it holds

$$f_{xy}(W) \geq f_{xy'}(W')$$

Maxima of Ψ are clearly Nash equilibria while, in general, the converse is not true. Considering that Ψ is defined on a finite set, a maximum, and thus a Nash equilibrium, always exists. Under a Nash equilibrium, a unit whose goal is to maximize its utility, has no advantage in moving their allocated data, under the standing assumption that only one data atom at a time can be moved. Notice that data atoms are to be interpreted as aggregations of data and the decision

of their size is part of the design of the algorithm. Clearly different levels of granularity will give rise to different game models including different Nash equilibria.

For any $W \in \mathcal{W}_p$, we define the Gibbs probability distribution over $\mathcal{X}^x(W)$ as

$$p_y(W, x) = \frac{e^{\gamma f_{xy}(W + e_{xy})}}{Z_\gamma}$$

where $\gamma > 0$ and where

$$Z_\gamma = \sum_{y \in \mathcal{X}^x(W)} e^{\gamma f_{xy}(W + e_{xy})}$$

is a normalizing factor.

A. The algorithm

The algorithm we are proposing below is a distributed and asynchronous algorithm where units activate at random independent times and either allocate or move their atoms, undertaking a relaxed stochastic version of the utility maximization.

The algorithm is mathematically described as a continuous time Markov process $Q_t : D_t \rightarrow \mathcal{X}$ on the set of partial allocations \mathcal{Q}_p . Precisely, units are assumed to be equipped with independent internal Poisson clocks with possibly different clicking rates. We denote by ν_x the clicking rate of unit x . When a unit activates it can either allocate a further data atom (if allocation is not completed yet) or move a data atom from one resource to another. The choice of the resource where either allocate or move the data atom is done according to the Gibbs probability: this is a classical choice in evolutionary game theory and will be amenable to a fairly complete theoretical analysis. The details of the algorithm are described below. We put $W(t) = W(Q_t)$.

- 1) Assume \bar{x} activates at time t . With probabilities

$$P_{\text{all}}(W(t), \bar{x}), \quad P_{\text{dis}}(W(t), \bar{x})$$

the resource \bar{x} will make, respectively, an allocation or a distribution move as explained below. Of course we are assuming that

$$P_{\text{all}}(W(t), \bar{x}) + P_{\text{dis}}(W(t), \bar{x}) = 1,$$

$$P_{\text{all}}(W(t), \bar{x}) = 0 \text{ if } W(t)^{\bar{x}} = \alpha_{\bar{x}}$$

$$P_{\text{dis}}(W(t), \bar{x}) = 0 \text{ if } W(t)^{\bar{x}} = 0$$

- 2) (ALLOCATION MOVE)

- Choose (\bar{x}, \bar{a}) uniformly at random in $(\mathcal{A} \setminus D_t) \cap \mathcal{A}_{\bar{x}}$
- Choose y^* according to the Gibbs probability $p_{y^*}(W(t), \bar{x})$
- Put $D_{t+} = D_t \cup \{(\bar{x}, \bar{a})\}$, and $Q_{t+} : D_{t+} \rightarrow \mathcal{X}$ by

$$Q_{t+}(x, a) = \begin{cases} y^* & \text{if } (x, a) = (\bar{x}, \bar{a}) \\ Q_t(x, a) & \text{if } (x, a) \in D_t \end{cases}$$

- 3) (DISTRIBUTION MOVE)


- Choose \bar{y} according to the probability $q_{\bar{y}} = W(t)_{\bar{x}\bar{y}} / W(t)^{\bar{x}}$.
- Choose (\bar{x}, \bar{a}) uniformly at random in $D_t \cap \mathcal{A}_{\bar{x}}$

- Choose y^* according to the Gibbs probability $p_{y^*}(W(t) - e_{\bar{x}\bar{y}}, \bar{x})$
- Put $D_{t+} = D_t$,

$$Q_{t+}(x, a) = \begin{cases} y^* & \text{if } (x, a) = (\bar{x}, \bar{a}) \\ Q_t(x, a) & \text{if } (x, a) \in D_t \setminus \{(\bar{x}, \bar{a})\} \end{cases}$$

The fact of using a noisy algorithm is crucial in our setting. In the following example we show a situation where a classical best response algorithm would remain stacked without completing the allocation.

Example 3: We are considering a line graph of four users as depicted below.



Each user x has $\alpha_x = \beta_x = 1$. Reliabilities are instead $\lambda_2 = 3$ while $\lambda_x = 1$ for $x = 1, 3, 4$. Assume we are in the partial allocation state $W \in \mathcal{W}_p$ given by

$$W := \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

Clearly, this partial allocation state could be reached by the group zero allocation state with positive probability (the key point is that unit 3 activates before unit 1 and chooses the most reliable resource). It is also clear that under a best response strategy ($\gamma = +\infty$) this state allocation is an equilibrium: unit 1 can not allocate, unless unit 3 moves its data to 4 but this will never happen because $f_{34}(W - e^{32}) < f_{32}(W - e^{32})$.

B. Theoretical results

In this section we analyze the behavior of the algorithm introduced above. We will essentially show two results. First, we prove that if the set of allocation \mathcal{Q} is not empty (i.e. condition (1) is satisfied), the algorithm above will find one in bounded time with probability 1. Second, we will show that, under a slightly stronger assumption than (1), in the double limit $t \rightarrow +\infty$ and then $\gamma \rightarrow +\infty$, the Markov process induced by the algorithm will always converge, in law, to a Nash equilibrium which is a global maximum of the potential function Ψ .

In order to prove such results, it will be necessary to go through a number of intermediate technical steps. First of all, it will be convenient to work directly with the process $W(t) = W(Q(t))$ which is also Markovian because of the way the transitions have been defined, considering that the results we are claiming, can all be expressed and established at this simpler level.

Given $W, W' \in \mathcal{W}_p$, we denote by $P_{W, W'}$ the transition probability from W to W' of the Markov chain underlying the process $W(t)$. \mathcal{L}_p denotes the graph on \mathcal{W}_p where an edge (W, W') is present if and only if $P_{W, W'} > 0$.

Our strategy will be to show that from any element $W \in \mathcal{W}_p$ there is a path in \mathcal{L}_p to some element $W' \in \mathcal{W}$. This, by standard Markov chain arguments, leads to the result that allocation will be achieved in bounded time with

probability 1. After, we will show that \mathcal{L}_p restricted to the set of allocations \mathcal{W} is irreducible and this will then yield the asymptotic result.

First we consider the problem of finding a path to an allocation state. Given $W \in \mathcal{W}_p$ we define the following subsets of units

$$\mathcal{X}^f(W) := \{x \in \mathcal{X} \mid W^x = \alpha_x\}$$

$$\mathcal{X}^{sat}(W) := \{x \in \mathcal{X} \setminus \mathcal{X}^f(W) \mid \exists y \in N_x \text{ s.t. } W(Q)_{y_i} < \beta_{y_i}\}$$

Units in $\mathcal{X}^f(W)$ are called fully allocated: these units have completed the allocation of their data under the state W . Units in $\mathcal{X}^{sat}(W)$ are called saturated: they have not yet completed their allocation, however, under the current state W , they can not make any action, neither allocate, nor distribute. Finally, define

$$\mathcal{W}_p^{sat} := \{W \in \mathcal{W}_p \setminus \mathcal{W} \mid \mathcal{X} = \mathcal{X}^f(W) \cup \mathcal{X}^{sat}(W)\}$$

It is clear that from any $W \in \mathcal{W}_p \setminus \mathcal{W}_p^{sat}$, some allocation move can be performed. Instead, if we are in a state $W \in \mathcal{W}_p^{sat}$, only possibly fully allocated units can make a distribution move. Notice that, because of condition (1), for sure there exist resources y such that $W_y < \beta_y$ and these resources are indeed exclusively connected to fully allocated units. The key point is to show that in a finite number of distribution moves it is always possible to move some data atoms from resources connected to saturated units to resources with available space: this will then make possible a further allocation move.

For any fixed $W \in \mathcal{W}_p$, we can consider the following graph structure on \mathcal{X} thought as set of resources: $\mathcal{H}_W = (\mathcal{X}, \mathcal{E}_W)$. Given $y_1, y_2 \in \mathcal{X}$, there is an edge from y_1 to y_2 if and only if there exists $x \in \mathcal{X}$ for which

$$W_{xy_1} > 0, \quad (x, y_2) \in \mathcal{E}$$

The edge from y_1 to y_2 will be indicated with the symbol $y_1 \rightarrow_x y_2$ (to also recall the unit x involved). The presence of the edge means that the two resources y_1 and y_2 are in the neighborhood of a common unit x which is using y_1 under W . This indicates that x can in principle move some of its data currently stored in y_1 into resource y_2 if this last one is available. We have the following technical result

Lemma 3: Suppose $(\mathcal{G}, \alpha, \beta)$ satisfies (1). Fix $W \in \mathcal{W}_p$ and let $\bar{y} \in \mathcal{X}$ be such that there exists $\bar{x} \in N_{\bar{y}}$ with $W^{\bar{x}} < \alpha_{\bar{x}}$. Then, there exists a sequence

$$\bar{y} = y_0, x_0, y_1, \dots, y_{t-1}, x_{t-1}, y_t \quad (8)$$

satisfying the following conditions

- (Sa) Both families of the y_k 's and of the x_k 's are each made of distinct elements;
- (Sb) $y_k \rightarrow_{x_k} y_{k+1}$ for every $k = 0, \dots, t-1$;
- (Sc) $W_{y_k} = \beta_{y_k}$ for every $k = 0, \dots, t-1$, and $W(Q)_{y_t} < \beta_{y_t}$.

Proof: Let $\mathcal{Y} \subseteq \mathcal{X}$ be the subset of nodes which can be reached from \bar{y} in \mathcal{H}_W . Preliminarily, we prove that there exists $y' \in \mathcal{Y}$ such that $W_{y'} < \beta_{y'}$. Let

$$\mathcal{Z} := \{x \in \mathcal{X} \mid \exists y \in \mathcal{Y}, W_{xy} > 0\}$$

and notice that, by the way \mathcal{Y} and \mathcal{Z} have been defined,

$$x \in \mathcal{Z}, (x, y) \in \mathcal{E} \Rightarrow y \in \mathcal{Y} \quad (9)$$

Suppose now that, contrarily to the thesis, $W_y \geq \beta_y$ for all $y \in \mathcal{Y}$. Then,

$$\begin{aligned} \sum_{x \in \mathcal{Z}} \alpha_x &\leq \sum_{y \in \mathcal{Y}} \beta_y \\ &= \sum_{y \in \mathcal{Y}} W_y \\ &= \sum_{y \in \mathcal{Y}} \sum_{x \in \mathcal{Z}} W_{xy} \quad (10) \\ &= \sum_{x \in \mathcal{Z}} W^x \\ &< \sum_{x \in \mathcal{Z}} \alpha_x \end{aligned}$$

where the first inequality follows from (9) and (1), the first equality from the contradiction hypothesis, the second equality from the definition of \mathcal{Z} , the third equality again from (9) and, finally, last inequality from the existence of \bar{x} . This is clearly absurd and thus proves our claim.

Consider now a path of minimal length from \bar{y} to \mathcal{Y} in \mathcal{H}_W :

$$\bar{y} = y_0 \rightarrow_{x_0} y_1 \rightarrow \dots \rightarrow y_{t-1} \rightarrow_{x_{t-1}} y_t$$

and notice that the sequence $\bar{y} = y_0, x_0, y_1, \dots, y_{t-1}, x_{t-1}, y_t$ will automatically satisfy properties (Sa) to (Sc). ■

We are now ready to prove the first main result.

Theorem 4: Assume that

- 1) $v_x > 0$ for every $x \in \mathcal{X}$ such that $\alpha_x > 0$,
- 2) $P_{\text{all}}(W, \bar{x}) > 0$ if $W^{\bar{x}} < \alpha_{\bar{x}}$,
- 3) $(\mathcal{G}, \alpha, \beta)$ satisfies (1).

Then, with probability 1, the Markov process $W(t)$ will be, after a finite number of jumps, in the set of allocations \mathcal{W} .

Proof: In order to prove the claim, it will be sufficient to show that from any $W \in \mathcal{W}_p$ there is a path in \mathcal{L}_p (the graph underlying the possible transitions of the process $W(t)$) to some element $W' \in \mathcal{W}$. We will prove it by a double induction process. To this aim we consider two indices associated to any $W \in \mathcal{W}_p \setminus \mathcal{W}$. The first one is defined by

$$m_W = \sum_{x \in \mathcal{X}} (\alpha_x - W^x) \geq 1$$

To define the second, consider any $\bar{x} \in \mathcal{X} \setminus \mathcal{X}^f(W)$. We can apply Lemma 3 to W and any $\bar{y} \in N_{\bar{x}}$ and obtain that we can find a sequence of agents $\bar{y} = y_0, x_0, y_1, \dots, y_{t-1}, x_{t-1}, y_t$ satisfying the properties (Sa), (Sb), and (Sc) above. Among all the possible choices of $\bar{x} \in \mathcal{X}$, $\bar{y} \in N_{\bar{x}}$ and of the corresponding sequence, assume we have chosen the one minimizing t and denote such minimal t by t_W . The induction process will be performed with respect to the lexicographic order induced by the pair (m_W, t_W) .

In the case when $t_W = 0$, it means we can find $\bar{x} \in \mathcal{X}$, $\bar{y} \in N_{\bar{x}}$ such that $W_{\bar{y}} < \beta_{\bar{y}}$. Therefore, under the allocation state W , the unit \bar{x} can allocate a further data atom to \bar{y} . Considering that the activation of the unit \bar{x} has positive probability because of assumptions 1. and 2., this shows that

W is connected to a W' such that $m_{W'} < m_W$. In case $m_W = 1$, this means that $W' \in \mathcal{W}$.

Consider now any $W \in \mathcal{W}_p \setminus \mathcal{W}$ such that $t = t_W > 1$. Let $\bar{x} \in \mathcal{X}$, $\bar{y} \in N_{\bar{x}}$ and the sequence $\bar{y} = y_0, x_0, y_1, \dots, y_{t-1}, x_{t-1}, y_t$ satisfying the properties (Sa), (Sb), and (Sc) above. In the allocation state W , the unit x_{t-1} , if activated (and this again has positive probability to happen because of assumptions 1. and 2.), can thus move an atomic piece of data from y_{t-1} to y_t . The new allocation state is $W' = W - e_{x_{t-1}y_{t-1}} + e_{x_{t-1}y_t}$. Since $W'_{y_{t-1}} < \beta_{y_{t-1}}$, for sure $t_{W'} < t_W$. The induction argument is thus complete. ■

We are now left with studying the Markov process $W(t)$ on \mathcal{W} . We start with the following

Proposition 5: $W(t)$, restricted to \mathcal{W} , is a time reversible Markov process. More precisely, it holds

$$\binom{\alpha}{W} e^{\gamma^{\Psi(W)} P_{W,W'}} = \binom{\alpha}{W'} e^{\gamma^{\Psi(W')} P_{W',W}} \quad \forall W, W' \in \mathcal{W} \quad (11)$$

where

$$\binom{\alpha}{W} := \frac{\prod_x \alpha_x!}{\prod_{x,y} W_{xy}!}$$

Proof: Notice that the only cases when $P_{W,W'}$ and $P_{W',W}$ are not both equal to 0 is for those pairs $W, W' \in \mathcal{W}$ such that $W' = W - e_{x\bar{y}} + e_{x\bar{y}'}$ for some $x, \bar{y}, \bar{y}' \in \mathcal{X}$. Assume this to be the case. Then, it follows from the way the distribution moves of the algorithm have been defined that

$$P_{W,W'} = v_x \frac{W_{x\bar{y}}}{\alpha_x} e^{\gamma_{f_{x\bar{y}'}}(W')}, \quad P_{W',W} = v_x \frac{W'_{x\bar{y}'}}{\alpha_x} e^{\gamma_{f_{x\bar{y}}}(W)}$$

Substituting in (11) and using relation (7), it is immediate to check that equality holds. ■

We now show that under a slight stronger assumption than (1), namely,

$$\sum_{x \in A} \alpha_x < \sum_{y \in N(A)} \beta_y \quad \forall A \subseteq \mathcal{X}, \quad (12)$$

the process $W(t)$ restricted to \mathcal{W} is ergodic. Denote by \mathcal{L} the graph \mathcal{L}_p restricted to the set \mathcal{W} . Notice that, as a consequence of time-reversibility, \mathcal{L} is an undirected graph. Ergodicity is equivalent to proving that \mathcal{L} is connected. We start with a lemma analogous to previous Lemma 3.

Lemma 6: Suppose $(\mathcal{G}, \alpha, \beta)$ satisfies (12) and let $W \in \mathcal{W}$. Then, for every $\bar{y} \in \mathcal{X}$, there exists a sequence (8) satisfying the conditions (Sa), (Sb), and (Sc) as in Lemma 3.

Proof: It is sufficient to follow the steps of the proof of Lemma 3 noticing that in (10) the first equality is now a strict inequality, while the last strict inequality becomes an equality. ■

If $W, W' \in \mathcal{W}$ are connected through a path in \mathcal{L} , we write that $W \sim W'$. Introduce the following distance on \mathcal{W} : if $W^1, W^2 \in \mathcal{W}$

$$\delta(W^1, W^2) = \sum_{x,y} |W_{xy}^1 - W_{xy}^2|$$

A pair $(W^1, W^2) \in \mathcal{W}$ is said to be minimal if

$$\delta(W^1, W^2) \leq \delta(W^{1'}, W^{2'}) \quad \forall W^{1'} \sim W^1, \forall W^{2'} \sim W^2$$

Notice that \mathcal{L} is connected if and only if for any minimal pair (W^1, W^2) , it holds $W^1 = W^2$.

Lemma 7: Let (W^1, W^2) be a minimal pair. Suppose $y \in \mathcal{X}$ is such that $W_y^1 < \beta_y$. Then, $W_{xy}^1 = W_{xy}^2$ for all $x \in \mathcal{X}$.

Proof: Suppose by contradiction that $W_{xy}^1 < W_{xy}^2$ for some $x \in \mathcal{X}$. Then, necessarily, there exists $y' \neq y$ such that $W_{xy'}^1 > W_{xy'}^2$. Consider then $W^{1'} = W^1 - e^{xy'} + e^{xy}$. Clearly, $\delta(W^{1'}, W^2) < \delta(W^1, W^2)$ and this contradicts the minimality assumption. Thus $W_{xy}^1 \geq W_{xy}^2$ for all $x \in \mathcal{X}$. This yields $W_y^2 < \beta_y$. Exchanging the role of W^1 and W^2 we obtain the thesis. ■

Proposition 8: If condition (12) holds true, the graph \mathcal{L} is connected.

Proof: Let (W^1, W^2) be any minimal pair. We will prove that W^1 and W^2 are necessarily identical. Consider any resource y . It follows from Lemma 6 that we can find a sequence $y = y_0, x_0, y_1, \dots, y_{t-1}, x_{t-1}, y_t$ satisfying the same (Sa), (Sb), and (Sc) with respect to the state allocation W^1 . Among all the possible sequences, choose one with t minimal for given y . We will prove by induction on t that $W_{xy}^1 = W_{xy}^2$ for all $x \in \mathcal{X}$.

If $t = 0$, it means that $W_y^1 < \beta_y$. It then follows from Lemma 7 that $W_{xy}^1 = W_{xy}^2$ for all $x \in \mathcal{X}$. Suppose now that the claim has been proven for all minimal pairs (W^1, W^2) and any $y \in \mathcal{X}$ for which $t < \bar{t}$ (w.r. to W^1) and assume that $y = y_0, x_0, y_1, \dots, y_{\bar{t}-1}, x_{\bar{t}-1}, y_{\bar{t}}$ satisfies the properties (Sa), (Sb), and (Sc) with respect to W^1 . Notice that the unit $x_{\bar{t}-1}$ can move a data atom from resource $y_{\bar{t}-1}$ into resource $y_{\bar{t}}$ under the state allocation W^1 and obtain $W^{1'} = W^1 - e^{x_{\bar{t}-1}y_{\bar{t}-1}} + e^{x_{\bar{t}-1}y_{\bar{t}}}$. Consider now W^2 and notice that Lemma 7 yields $W_{x_{\bar{t}-1}y_{\bar{t}}}^2 = W_{x_{\bar{t}-1}y_{\bar{t}}}^1 < \beta_{y_{\bar{t}}}$. Define

$$W^{2'} = \begin{cases} W^2 & \text{if } W_{x_{\bar{t}-1}y_{\bar{t}-1}}^2 = 0 \\ W^2 - e^{x_{\bar{t}-1}y_{\bar{t}-1}} + e^{x_{\bar{t}-1}y_{\bar{t}}} & \text{if } W_{x_{\bar{t}-1}y_{\bar{t}-1}}^2 > 0 \end{cases}$$

Clearly, $\delta(W^{1'}, W^{2'}) \leq \delta(W^1, W^2)$ and this implies that also $W^{1'}, W^{2'}$ is a minimal pair. Notice that $y = y_0, x_0, y_1, \dots, y_{\bar{t}-1}$ satisfies (Sa), (Sb), and (Sc) with respect to $W^{1'}$. Therefore, by the induction hypotheses, it follows that $W_{xy}^{1'} = W_{xy}^{2'}$ for all $x \in \mathcal{X}$. Since $W_{xy}^1 = W_{xy}^{1'}$ and $W_{xy}^2 = W_{xy}^{2'}$, result follows immediately. ■

Propositions 5 and 8 immediately yield the following final result.

Corollary 9: Assume that (12) holds true. Then $W(t)$, restricted to \mathcal{W} , is an ergodic time reversible Markov process whose unique invariant probability measure is given by

$$\mu_\gamma(W) = \frac{\left(\frac{\alpha}{W}\right) e^{\gamma\Psi(W)}}{\sum_{W' \in \mathcal{W}} \left(\frac{\alpha}{W'}\right) e^{\gamma\Psi(W')}}.$$

Remark: Notice that when $\gamma \rightarrow +\infty$, the invariant probability μ_γ converges to the probability μ_∞ concentrated on

the set $\operatorname{argmax}_{W \in \mathcal{W}} \Psi(W)$ of state allocations maximizing the potential and given by, for $W \in \operatorname{argmax}_{W \in \mathcal{W}} \Psi(W)$,

$$\mu_\infty(W) = \frac{\left(\frac{\alpha}{W}\right)}{\sum_{W' \in \mathcal{W}} \left(\frac{\alpha}{W'}\right)}$$

Thus, if γ is small, the distribution of the process $W(t)$ for t sufficiently large will be close to Nash equilibria.

In this paper we will assume that $v_x = v\alpha_x$ for some $v > 0$, namely that units activation rates is proportional to the amount of data they need to back up. We will consider two possibilities for the allocation and distribution probabilities $P_{\text{all}}(W(t), x)$ and $P_{\text{dis}}(W(t), x)$:

$$\begin{aligned} P_{\text{all}}(W(t), x) &= \frac{\alpha_x - W(t)^x}{\alpha_x}, \\ P_{\text{all}}(W(t), x) &= \begin{cases} 1 & \text{if } W(t)^x < \alpha_x \\ 0 & \text{otherwise} \end{cases} \\ P_{\text{dis}}(W(t), x) &= \frac{W(t)^x}{\alpha_x}, \\ P_{\text{dis}}(W(t), x) &= \begin{cases} 0 & \text{if } W(t)^x < \alpha_x \\ 1 & \text{otherwise} \end{cases} \end{aligned}$$

In the first case, the probability of an allocation move is proportional to the amount yet to be allocated while in the second case, no distribution move takes place before full allocation is reached.

III. SIMULATION

In this section we present a number of numerical simulations that validate the theoretical results and show the performance of the algorithm in terms of various parameters describing the speed of convergence, resources congestions, global utility, and complexity of the interconnections.

For the sake of readability we gather below the standing assumptions and parameters we have used.

- The number of units is denoted by n and assumed to be even. Most of our simulations are for $n = 50$ but we have also studied scalability issues by considering $n = 100$ and $n = 1000$.
- We have considered two possible interconnection topologies: the complete graph and a random regular graph of degree 10.
- Time has been assumed to be discrete, assuming that at every time instant a unit x is chosen with a probability v_x proportional to α_x . Moreover, allocation and distribution moves are chosen according to the probabilities:

$$P_{\text{all}}(W(t), x) = \begin{cases} 1 & \text{if } W(t)^x < \alpha_x \\ 0 & \text{otherwise} \end{cases}$$

- Time horizon has been fixed to be $T = 2\sum \alpha_x$ so that we have allowed up to two moves per data atom (one allocation and one possible distribution). It turns out that in all experiments carried on such time horizon has been sufficient for completing the allocation of all data atoms and also getting very close to a Nash equilibrium. Denote by W^∞ the final allocation state of the system after time T has elapsed.

- The parameter γ appearing in the Gibbs distribution have been chosen to be time-varying with

$$\gamma(t) = \gamma(t-1) + \frac{1}{\lambda_{max} * 100}$$

where λ_{max} is the maximum reliability of the resources. This is a typical choice done on such best response dynamics, even if the theoretical result expressed below can not directly be applied to insure convergence. In real world applications, where adaptation to a time-varying scenario (e.g. addition or deletion of units, change in the topology) is needed, γ must be instead kept bounded.

- We have assumed units to have the same free space to offer $\beta = 50$ and to have possibly different amounts α_x to allocate.
- We have assumed units to split into two subsets of equal size \mathcal{X}_1 and \mathcal{X}_2 characterized by two different reliability levels, respectively, $\lambda_1 = 0.5$ and $\lambda_2 = 0.8$.
- The congestion parameter is chosen to be $k_c = 1$ while we will consider different values for k_a .

Moving data from one resource to another one is an expensive task which must be carefully monitored in real applications. To this aim, we have introduced the index v_{moves} which computes the number of allocation or distribution moves per piece of data throughout the dynamics. In formula, if m_i is the total number of moves performed by agent i during the run of the algorithm, we put

$$v_{moves} = \frac{1}{n} \sum_{x \in \mathcal{X}} \frac{m_i}{\alpha_i}.$$

The global utility of the system in the allocation state W is defined as $F(W) = \sum_{x,y \in \mathcal{X}} W_{xy} f_{xy}(W)$. Put $F^* = \max_{W \in \mathcal{W}} F(W)$.

$$\rho := \frac{F(W^\infty)}{F^*}$$

measures how the solution found is performing with regards to the global utility (notice that the potential Ψ does not coincide with F so that W^∞ is not a-priori a maximum of F).

Units give preference to the most trusted resource in \mathcal{X}_2 , however, depending on the amount of data they need to allocate and on the type of resources in their neighborhood, they need to use also less trusted resources in \mathcal{X}_1 . We define the mean and the variance of the satisfaction level as

$$\bar{\Lambda} := \frac{1}{n} \sum_{x \in \mathcal{X}} \sum_{y \in N_x} \frac{W_{xy}^\infty}{\alpha_x} \lambda_y,$$

$$\Lambda_{var} := \frac{1}{n} \sum_{x \in \mathcal{X}} \left(\sum_{y \in \mathcal{X}} \frac{W_{xy}^\infty}{\alpha_x} \lambda_y - \bar{\Lambda} \right)^2$$

If λ_1 and λ_2 are taken to be the probability that if contacted at a random time the resource is available to give access to the stored data, $\bar{\Lambda}$ can be interpreted as the probability that a piece of data can be recovered when requested at some random time.

The presence of the congestion term in the utility function should insure that all resources with the same λ should in

principle be used equally. We measure the mean and variance of the congestion level of resources in \mathcal{X}_i by

$$\bar{C}^i := \frac{1}{n\beta} \sum_{y \in \mathcal{X}_i} W_y^\infty,$$

$$C_{var}^i := \frac{1}{n} \sum_{y \in \mathcal{X}_i} (\beta^{-1} W_y^\infty - \bar{C})^2$$

Finally, we consider the in and out mean degrees measuring the topological complexity of the subgraph consisting of the edges (x,y) for which $W_{xy}^\infty > 0$:

$$d^+ := \frac{1}{n} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{X}} \mathbb{1}_{\{W_{xy}^\infty > 0\}},$$

$$d_i^- := \frac{2}{n} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{X}_i} \mathbb{1}_{\{W_{xy}^\infty > 0\}}, \quad i = 1, 2$$

Below we present four sets of examples. The first two are with $n = 50$ units, α_x constant, and different topologies. The third one deals instead with a situation where units have different α_x . Finally, the fourth example deals with a larger number of users.

Example 4: Consider a case with $n = 50$ users and $\alpha_x = \alpha = 45$ for every unit x . It follows from Proposition 2 that allocation is possible in any regular graph. Also notice that Theorem 4 can be applied. We first assume that the graph is complete and then we analyze the case of a regular graph of degree 10.

The following table shows the value of the various performance parameters varying the aggregation term k_a .

TABLE I: Complete Graph

| | $k_a = 0$ | $k_a = 0.25$ | $k_a = 0.45$ |
|-------------------|--------------------|--------------|--------------|
| v_{moves} | 1.6271 | 1.3068 | 1.2548 |
| $\bar{\Lambda}$ | 0.6667 | 0.6592 | 0.6593 |
| Λ_{var} | $6.4818 * 10^{-4}$ | 0.0119 | 0.0122 |
| \bar{C}^1 | 0.8000 | 0.8450 | 0.8442 |
| \bar{C}_{var}^1 | $9.5680 * 10^{-4}$ | 0.1149 | 0.1195 |
| \bar{C}^2 | 1 | 0.9550 | 0.9558 |
| \bar{C}_{var}^2 | 0 | 0.0280 | 0.0261 |
| d^+ | 44.8460 | 9.5420 | 9.6720 |
| d_1^- | 43.9280 | 9.1720 | 9.1280 |
| d_2^- | 45.7640 | 9.9120 | 10.2160 |
| ρ | 0.9787 | 0.6812 | 0.6796 |

The fact that, on average, each user makes less than two moves per each atom, suggests that at the end of the allocation the system is already near to the equilibrium.

Notice the effect of the aggregation term on this parameter: when $k_a = 0$, the parameter v_{moves} is considerably higher as agents tend to move their data to leverage the congestion level over all the available resources. When the aggregation term is present this phenomenon is much reduced as agents have an incentive to keep data atoms together. The congestion indices confirm that agents preferably allocate over more trusted resources and that resources are used sufficiently equally. We can also see that the aggregation term brings to a slightly greater usage of worse resources. The parameters

where the influence of the aggregation term is even clearer, are the degrees: in the first case almost every edge is used while in the second two cases the number is sensibly lower. Finally, concerning the utility parameter ratio ρ , notice that our algorithm performs very well for $k_a = 0$ while we witness a certain degradation when $k_a > 0$.

Example 5: Consider now the same case with 50 users where the underlying graph is regular with degree 10. As in the previous case, the table shows the parameters in the case with $k_a = 0$, $k_a = 0.25$ and $k_a = 0.45$.

TABLE II: Regular Degree 10

| | $k_a = 0$ | $k_a = 0.25$ | $k_a = 0.45$ |
|-----------------|-----------|--------------|--------------|
| v_{moves} | 1.4187 | 1.2185 | 1.1714 |
| $\bar{\Lambda}$ | 0.6667 | 0.6596 | 0.6606 |
| Λ_{var} | 0.0019 | 0.0136 | 0.0143 |
| \bar{C}^1 | 0.8000 | 0.8422 | 0.8364 |
| C_{var}^1 | 0.0011 | 0.1214 | 0.1350 |
| \bar{C}^2 | 1 | 0.9578 | 0.9636 |
| C_{var}^2 | 0 | 0.0261 | 0.0251 |
| d^+ | 9.9560 | 6.2580 | 6.3700 |
| d_1^- | 9.9240 | 5.9400 | 6.2520 |
| d_2^- | 9.9880 | 6.5760 | 6.4880 |
| ρ | 0.9784 | 0.8872 | 0.9297 |

These data confirm the good performance properties of the algorithm even when the graph of connection is of considerably less complexity. Notice how, when the aggregation term is present, the number of resources used per agent is something above 6 to be compared to the previous case where it was almost 10. In other terms, imposing a constrained communication pattern, it does not degrade performance and helps keeping complexity at a lower level.

Figures 1 and 2 show the difference between the graph we impose and the graph with edges that are used by the users in the case with $k_a = 0.25$.




Fig. 1: Underlying Network

Next example shows that variability of the amount of the data to be allocated by the various units does not significantly change the overall performance of the algorithm.

Example 6: In this example we consider the same number of $n = 50$ users split into five equal subfamilies with varying $\alpha = 35, 40, 45, 50, 55$ (so that the average is still 45). The underlying graph is assumed to be random regular with




Fig. 2: Used Edges

degree 10 and $k_a = 0.45$. A-priori there is no guarantee that an allocation exists in this case, however, the algorithm always finds one. The table below reports the value of the various performance indices which turn out to be quite close to the case of constant α_x .

TABLE III: Different α

| | |
|-----------------|--------|
| v_{moves} | 1.1552 |
| $\bar{\Lambda}$ | 0.6613 |
| Λ_{var} | 0.0138 |
| \bar{C}^1 | 0.8387 |
| C_{var}^1 | 0.1464 |
| \bar{C}^2 | 0.9613 |
| C_{var}^2 | 0.0328 |
| d^+ | 6.4040 |
| d_1^- | 6.1200 |
| d_2^- | 6.6880 |

Finally, next example shows that the algorithm has good scalability properties.

Example 7: In this example we consider and compare three communities with $n = 50, 100, 1000$ always connected through a random regular graph of degree 10. We assume that $\alpha_x = 45$ for all units. The aggregation parameter is $k_a = 0.45$.

TABLE IV: Asymptotic Behavior

| | 50 | 100 | 1000 |
|-----------------|--------|--------|--------|
| v_{moves} | 1.1714 | 1.1490 | 1.1304 |
| $\bar{\Lambda}$ | 0.6606 | 0.6605 | 0.6566 |
| Λ_{var} | 0.0143 | 0.0146 | 0.0114 |
| \bar{C}^1 | 0.8364 | 0.8370 | 0.8604 |
| C_{var}^1 | 0.1350 | 0.1262 | 0.1068 |
| \bar{C}^2 | 0.9636 | 0.9630 | 0.9396 |
| C_{var}^2 | 0.0251 | 0.0183 | 0.0616 |
| d^+ | 6.3700 | 6.2840 | 6.1902 |
| d_1^- | 6.2520 | 5.9380 | 6.0004 |
| d_2^- | 6.4880 | 6.6300 | 6.3800 |

IV. CONCLUSIONS

In this paper we have proposed, in a game theoretic framework, a peer-to-peer decentralized storage model where a network of units are, at the same time, end users needing to allocate externally a back up of their data, as well

storage resources for other users. We have proposed a novel fully distributed algorithm where units, connected through a network, activate autonomously at random time and either allocate or move pieces of their data among the neighboring available resources. Actions taken by the units are noisy best response actions with respect to utility functions which incorporate the congestion of the resources, their reliability, as well as possibly, the fragmentation of the stored data. The algorithm has been claimed to converge, with probability 1, to a Nash equilibrium of the game and several numerical simulations have here validated this claim. In a forthcoming paper, we will propose a detailed theoretical analysis of our algorithm.

We believe that there are several challenges related to the peer-to-peer storage model which have not yet been satisfactorily addressed by pure mathematical model. Some of them are reported below and will be the goals of our future research.

- In realistic scenarios, for security reasons, units need to allocate more than one copy of their own data. This poses new issues and constraints as it becomes fundamental that copies of the same data are not stored in the same resource. Most of our theoretical results, including the matching problem equivalence, will need to be completely revisited to be applied in this new scenario.
- Our simulations show that the algorithm has very good convergence properties. It will be of interest to establish theoretical bounds on the convergence time.
- In our model, units, thought as resources, are assumed to be always on and available, if not yet fully congested, for storage actions. It would be of interest to let the possibility that resources might be off at certain times and to connect this behavior to the reliability parameter present in the utility functions.
- Units, in our model, are completely anonymous and resources do not make any filter on new allocation requests. More interesting models should incorporate trust formation mechanisms where more trusted units (when thought as resources) should have a vantage in finding place to store their data.

ACKNOWLEDGMENT

We acknowledge that this work has been done while Barbara Franci was a PhD Student sponsored by a Telecom Italia grant.

REFERENCES

- [1] V. Jalaparti, G. Nguyen, I. Gupta, M. Caesar, "Cloud Resource Allocation Games", Technical Report, University of Illinois. <http://hdl.handle.net/2142/17427>
- [2] R. W. Rosenthal, "A Class of game Possessing Pure-Strategy Nash Equilibrium", Int. J. Game Theory 2, 6567.
- [3] T. Roughgarden, E. Tardos, "How Bad is Selfish Routing?", Journal of ACM, vol. 49, no. 2, pp. 236259, 2002.
- [4] C. Tekin, M. Liu, R. Southwell, J. Huang, S.H. A. Ahmad, "Atomic Congestion Game on Graphs and its Application in Networking", IEEE/ACM Transaction on Networking, vol. 20, no. 5, pp. 1541-1552, 2012
- [5] H. Tembine, E. Altman, R. El-Azouzi, Y. Hayel, "Evolutionary Games in Wireless Networks", IEEE Trans. on Systems, Man, and Cybernetics, Part B, Cybernetics, vol. 40, pp. 634-646, 2010.
- [6] W.H. Sandholm, "Population Games and Evolutionary Dynamics" Cambridge, MA: The MIT press, 2010.